

IN THE CLAIMS

Please cancel claim 1 without prejudice or disclaimer.

Please amend claims 2, 3, 4, 5, 6, 7 and 8 as indicated below.

This listing of claims will replace all prior versions, and listings, of claims in the application.

Listing of Claims:

Claim 1 (cancelled)

Claims 2 (currently amended) The method as described in claim [[1]] 5 wherein the entries with single value attributes are stored in the merged table.

Claim 3 (currently amended) The method as described in claim [[1]] 5 wherein the entries with multiple value attributes are stored in the overflow table.

Claim 4 (currently amended) The method as described in claim [[1]] 5 wherein the overflow table is an attribute table.

Claim 5 (currently amended) ~~The method as described in claim 1~~ A method for storing data that has at least some entries with multiple value attributes, comprising the steps of:

profiling the data to determine whether the data should be stored in an attribute table or, alternatively, in a merged table and an overflow table; and
storing the data optimally based on the profiling step;

wherein a majority of the data is stored in the merged table and a set of additional values for the multiple value attributes are stored in the overflow table.

Claim 6 (currently amended) The method as described in claim [[1]] 5 wherein the profiling step parses the data to identify entries with single value attributes.

Claim 7 (currently amended) The method as described in claim [[1]] 5 wherein the profiling step parses the data to identify given operations that are performed on the data once stored.

Claim 8 (currently amended) The method as described in claim [[1]] 5 wherein the data is stored in a relational database backing store.

Claim 9 (withdrawn) A method for storing data that has at least some entries with multiple value attributes, comprising the steps of:

profiling the data to determine whether the data should be stored in per attribute tables or in a merged table and an overflow table;

storing the information optimally according to the profiling step; and

keeping pointers for future operations on the data in the merged and overflow tables.

Claim 10 (withdrawn) The method as described in claim 9 wherein a majority of the data is stored in the merged table and a small set of additional values for the multiple value attributes are stored in the overflow table.

Claim 11 (withdrawn) The method as described in claim 9 wherein the profiling step is performed as a function of a number of data entries having multiple value attributes.

Claim 12 (withdrawn) The method as described in claim 9 wherein the profiling step is performed by identifying operations that will be performed on the data once stored.

Claim 13 (withdrawn) The method as described in claim 9 wherein the data is stored in a relational database backing store.

Claim 14 (withdrawn) The method as described in claim 9 wherein the merged table and the overflow table are stored in the relational database backing store.

Claim 15 (withdrawn) A database schema for a directory service having a backing store, comprising:

a first table for storing data for entries having single value attributes; and
a second table for storing data for entries having multiple value attributes;
wherein a majority of the entries have single value attributes and a small set of the entries have multiple value attributes.

Claim 16 (withdrawn) The database schema as described in claim 15 wherein the second table is a per attribute table.

Claim 17 (withdrawn) The database schema as described in claim 15 wherein the directory service is implemented according to a Lightweight Directory Access Protocol (LDAP).

Claim 18 (withdrawn) The database schema as described in claim 17 wherein the backing store is a relational database.

Claim 19 (withdrawn) A directory service, comprising:

a directory organized as a naming hierarchy having a plurality of entries each represented by a unique identifier;
a relational database management system having a backing store for storing directory data;
a database schema, comprising:
a first table for storing data for entries having single value attributes; and
a second table for storing data for entries having multiple value attributes;
wherein a majority of the entries have single value attributes and a small set of the entries have multiple value attributes.

Claim 20 (withdrawn) The directory service as described in claim 19 wherein the directory is compliant with the Lightweight Directory Access Protocol (LDAP).

Claim 21 (withdrawn) In a directory service having a directory organized as a naming hierarchy, the hierarchy including a plurality of entries each represented by a unique identifier, the improvement comprising:

 a relational database management system having a backing store for storing directory data according to a schema comprising:

 a first table for storing data for entries having single value attributes;
 and

 a second table for storing data for entries having multiple value attributes;

 wherein a majority of the entries have single value attributes and a small set of the entries have multiple value attributes.

Claim 22 (withdrawn) In the directory service as described in claim 21 wherein the directory is compliant with the Lightweight Directory Access Protocol (LDAP).

Claim 23 (withdrawn) A method for storing data that has at least some entries with multiple value attributes, comprising the steps of:

 storing data for entries having single value attributes in a first table;

 storing data for entries having multiple value attributes in a second table; and

 wherein a majority of the entries have single value attributes and a small set of the entries have multiple value attributes.